



Grid-like units help deep learning agent to navigate

Ken Cheng¹

Published online: 14 June 2018
© Psychonomic Society, Inc. 2018

Summary

An artificial-intelligence model based on deep learning developed units in a hidden layer that resembled mammalian grid cells in the hippocampus when the agent was taught to integrate paths. The full model performed sophisticated navigational tasks—in some cases even better than a human.

Keywords Navigation · Spatial learning · Modeling · Artificial intelligence

Neurophysiological work on mammalian place cells and grid cells in the hippocampal formation (see Cheng & Jeffery, 2017) won the Nobel Prize in Physiology or Medicine in 2014. Place cells, found in the hippocampus, fire the most at particular places in the environment in which the animal is traveling, irrespective of how the animal got there or which way it is facing. Grid cells, found in the entorhinal cortex, show increased firing in a grid of locations in space. The grid forms a triangular tessellation at a particular scale and orientation (Cheng & Jeffery, 2017). Grid cells are thought to play a role in path integration, in which the animal keeps track of distances and directions that it has traveled. Behavioral neuroscientists have been inspired to model the workings of the hippocampal system over the years. A recent venture featured in *Nature Magazine*, however, spanned beyond the confines of academia. In collaboration with University College London, a sizeable team from Google DeepMind, with lead author Andrea Banino, combined insights from the behavioral neuroscience of the mammalian hippocampal formation with deep learning in artificial intelligence to create navigating agents (Banino et al., 2018).

Deep learning is based on neural networks with multiple sequential layers of computational units (Savelli & Knierim, 2018). In a classification task, such a network would be trained on a set of cases, instances that do or do not belong to the category in

question. The network somehow comes up with rules to classify new instances. The fudging term “somehow” is needed here, because scientists usually have trouble figuring out how the network works; deep learning comes with a black-box flavor. Nevertheless, deep learning has achieved superhuman feats in a literal sense, such as a program called AlphaGo that can defeat the human champion in the demanding game of Go (see <https://deepmind.com/research/alphago/>).

The deep-learning navigational system of Banino et al. (2018) contains multiple components. Capitalizing on what is known about rodent navigation, the DeepMind–UCL team first used deep learning to train their protégé to integrate paths. Inputs of linear and angular velocity were fed to the agent as it walked various paths that rodents had traversed. Interestingly, units in a middle layer that worked like grid cells emerged from the training. These grid units—to distinguish them from biological cells—responded the most in a triangular grid of locations. Such grid units comprised a sizeable minority (~23%) of the units in the hidden layer. Other units resembled other types of spatially tuned cells, such as border cells, which fire the most near some border of space (see Cheng & Jeffery, 2017). Yet others were place-cell like, while many showed miscellaneous nondescript patterns of responding. Crucially, grid-like units only turned up with some noise injected into the system; the authors accomplished this by randomly dropping out a proportion of the units at each step. Like grid cells, the grid units came in multiple scales. Model comparison of the distribution of scales showed that a tri-modal distribution provided the best description, with the scales having a factor of ~1.5 between successive modes. The authors hinted that grid units might be helping path integration by giving a metric scale for space, but exactly how these units contribute to path integration remains unclear.

✉ Ken Cheng
ken.cheng@mq.edu.au

¹ Department of Biological Sciences, Macquarie University, Sydney, NSW 2109, Australia

Next, the path integrating agent was amplified with a visual input module at the front end and an executive actor-critic at the back end, and set loose on challenging navigational tasks. Importantly, the actor-critic executive took input from the layer containing the grid units. Rooms with obstacles and doors, a version of a sunburst maze, and mazes with multiple arms that provided a test of shortcutting abilities were foisted on the grid cell agent, as the authors named their star protégé, and other agents, including a human. One other agent, called the place cell prediction agent, also used the path-integration module, but took inputs from layers further downstream from the grid-unit-containing layer. And a garden-variety deep-learning system not capitalizing on the path-integration module at all was also in the mix.

The short story on the outcomes is that the grid-unit agent came out tops—yes, even beating the lightly trained human (although the human had far less task training than the artificial agents, the authors pointed out that the human would have had ample real-world training before coming to the task). On a range of tests, including a test on negotiating shortcuts in mazes, the grid-unit agent typically did better than other artificial agents by substantial margins, with effect sizes mostly exceeding 1.0, and sometimes much larger.

In commenting on this big collaborative effort in *Nature*, Savelli and Knierim (2018) found it interesting that the trained network developed grid-like units. From this outcome, the commentators suggested that “something special about grid cells’ activity patterns” (p. 314) helps rodents to integrate paths. But these authors also pointed out that it is difficult to pin down that something, and called for advances to make deep learning more understandable.

Vestibular signals, whose instantiation was used by Banino et al. (2018) to train their artificial pupil, are not the only signals that can be used for path integration. Insects use an external compass in the sky called the sky compass (Stone et al., 2017), and their path integration system has been modeled, the most recent to my knowledge being Stone et al.’s (2017) neurobiologically inspired model. Stone et al.’s model does not represent distance explicitly; the odometry (measure of distance traveled) is coded implicitly by the strength of a signal in a particular compass direction. Banino et al.’s (2018) system shows that distance coding is important for their

protégé’s success. An insect model taken as a whole, then, might not work as well at navigation. But Google DeepMind and others should consider hybrid mammal–insect systems for real-world navigation. The insect’s external compass works much better than the internal vestibular compass of rodents. The reason is that an external compass can be checked at every step anew. Each step, then, has its own independent source of error, and error in executing Step n is not added cumulatively to the errors from Steps 1 to $n - 1$. In contrast, a vestibular system piles errors up cumulatively, each step starting off from the already accumulated errors of all previous steps. This sensory difference is one big reason, perhaps *the* reason, why huge-brained humans could do worse than miniscule-brained desert ants in tasks of path integration.

Mixing an external compass system with DeepMind’s deep learning navigational system might make a superior navigator in the real outdoors world in which a sky compass is available, at least during daylight hours. Navigating robots using a sky compass were invented 2 decades ago (Lambrinos, Möller, Labhart, Pfeifer, & Wehner, 2000). Hybrid systems are worth exploring. After all, deep learning systems all require inputs of features that they can extract. Those in artificial intelligence say that deep learning is only as good as the inputs it gets; the better the input, the better the outcome (Michael Milford, personal communication, February 2018).

References

- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., ... Kumaran, D. (2018). Vector-based navigation using grid-like representations in artificial agents. *Nature*, *557*, 429–433. doi: <https://doi.org/10.1038/s41586-018-0102-6>
- Cheng, K., & Jeffery, K. J. (2017). Spatial cognition. In J. Call (Ed.), *APA handbook of comparative psychology* (Vol. 2, pp. 463–483). Washington, DC: American Psychological Association.
- Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., & Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, *30*, 39–64.
- Savelli, F., & Knierim, J. J. (2018). AI mimics brain codes for navigation. *Nature*, *557*, 313–314.
- Stone, T., Webb, B., Adden, A., Weddig, N. B., Honkanen, A., Templin, R., ... Heinze, S. (2017). An anatomically constrained model for path integration in the bee brain. *Current Biology*, *27*, 3069–3085.